

Truthfulness, Proportional Fairness, and Efficiency*

Richard Cole [†]

Vasilis Gkatzelis [‡]

Gagan Goel [§]

July 10, 2012

Abstract

How does one allocate a collection of resources to a set of *strategic* agents in a fair and efficient manner without using money? For in many scenarios it is not feasible to use money to compensate agents for otherwise unsatisfactory outcomes. This paper studies this question, looking at both fairness and efficiency measures.

- We employ the *proportionally fair* solution, which is a well-known fairness concept for money-free settings. But although finding a proportionally fair solution is computationally tractable, it cannot be implemented in a truthful fashion. Consequently, we seek approximate solutions. We give several truthful mechanisms which achieve proportional fairness in an approximate sense. We use a strong notion of approximation, requiring the mechanism to give each agent a good approximation of its proportionally fair utility. In particular, one of our mechanisms provides a better and better approximation factor as the minimum demand for every good increases. A motivating example is provided by the massive privatization auction in the Czech republic in the early 90s.
- With regard to efficiency, prior work has shown a lower bound of 0.5 on the approximation factor of any swap-dictatorial mechanism approximating a social welfare measure even for the two agents and multiple goods case. We surpass this lower bound by designing a non-swap-dictatorial mechanism for this case. Interestingly, the new mechanism builds on the notion of proportional fairness.

*This work was supported in part by NSF grant CCF0830516.

[†]cole@cims.nyu.edu. Courant Institute, New York University, N.Y.

[‡]gkatz@cims.nyu.edu. Courant Institute, New York University, N.Y.

[§]gagangoel@google.com. Google Research, New York, N.Y.

1 Introduction

How does one allocate a collection of resources to a set of *strategic* agents in a fair and efficient manner without using money? This is a fundamental problem with many applications since in many scenarios payments cannot be solicited from agents; for instance, different teams compete for a set of shared resources in a firm, and the firm cannot solicit payments from the teams to make allocation decisions. Despite the applicability of mechanism design without money, much of the work in this area relies on enforcing payments. Motivated by this, we study the above question in this paper. We look at both fairness and efficiency measures.

One practical issue that arises in mechanism design without money is to put the utilities of the agents on a common scale. For example, happiness could mean different things to different people and cannot be compared as such. When payments can be used, a standard approach is to measure the utilites in terms of *money*. In the absence of money, one way to overcome this difficulty is to look for *scale-free* solutions, i.e., if an agent scales her valuations up or down, the solution should remain the same. When maximizing social welfare (SW), one can achieve this by first normalizing the values of the agents so that they add up to a common number (1 say) and then maximizing the welfare with these normalized values. For the case of fairness, the *proportionally fair* solution is a well-known fairness concept that is scale-free. In brief, a proportionally fair (PF) solution is a Pareto optimal solution \mathcal{O} which compares favorably to any other Pareto optimal solution \mathcal{O}' in the sense that when switching from \mathcal{O}' to \mathcal{O} , percentage gains outweigh percentage losses. The PF solution was first proposed in the TCP literature and is used widely in many practical scenarios [12].

Maximizing social welfare with normalized values using truthful mechanisms was first studied by Guo and Conitzer [9]. They studied the mechanism design for the special case of two players and many items to better understand the structure of good mechanisms maximizing social welfare. Even in this special case, they showed that the problem is difficult by proving that no mechanism from a general class called *increasing-price* mechanisms that use artificial currency can yield better than 0.5 approximation; they left it as an open question to overcome this bound. Later on, Han et al. [10] showed the same negative result of 0.5 for an even more general class of mechanisms called *swap-dictatorial*¹ mechanisms. In this work, we break this bound of 0.5. Our main contribution here is to show a connection between the PF and SW allocations; then we exploit this connection to give an interesting non-swap-dictatorial mechanism which overcomes the 0.5 bound.

We then focus on the design of truthful mechanisms for achieving proportional fairness. Even for simple instances involving just two agents and two items, it is not difficult to show that no truthful mechanism can obtain a PF solution in general. Hence one has to look for approximate solutions. In this work, we ask for a strong notion of approximation requiring the mechanism to give each agent a good approximation of its true proportionally fair utility.

Also, we note that the PF solution is related to the market equilibrium prices when agents are assumed to each have a *unit budget* of some artificial currency. In this case, the PF allocation is the same as the allocation at market equilibrium prices, and is captured via the Eisenberg-Gale program [8, 7, 11]. This connection of the PF solution to market equilibrium prices is of independent interest for the design of truthful mechanisms in this setting.

A specific real world example of the use of artificial currency for achieving or approximating

¹This class contains all mechanisms that (randomly) choose one of the two bidders and allow her to choose her preferred bundle of items from a predefined set; the other bidder receives the remaining items. For further discussion on swap-dictatorial mechanisms and their importance in money-free mechanism design, see [9] and references therein.

fair outcomes, which motivates our final result, is that of the privatization auctions that took place in Czechoslovakia. In the early 90s, the Czech government sought to privatize the state owned firms dating from the then recently ended communist era. The government’s goal was two-fold — first to distribute shares of these companies to their citizens in a fair manner and to calculate the market prices of these companies so that the shares could be traded in the open market after the initial allocation. To this end, they ran an auction, as described in [1]. Citizens could choose to participate by buying 1000 vouchers at a cost of 1,000 Czech Crowns, about \$35, a fifth of the average monthly salary. Over 90% of those eligible participated. These vouchers were then used to bid for shares in the available 1,491 firms. We believe that the PF allocation provides a very appropriate solution for this example, both to calculate a fair allocation and to compute market prices. Our mechanism solves the problem of efficiently discovering approximately PF allocations in a truthful fashion for such natural scenarios where there is high demand for each resource.

Our results. We start in Section 3 with the goal of (approximately) maximizing the social welfare using truthful mechanisms. Previous work showed that no non-trivial approximation factors can be achieved by any truthful mechanism for unbounded numbers of items and bidders [10]. This work also showed that no swap-dictatorial mechanism can yield an approximation factor better than the trivial 0.5 for the two-bidder case. Accordingly, the authors asked whether a truthful mechanism breaking this bound of 0.5 exists. We give a positive answer by designing a truthful non-swap-dictatorial mechanism closely related to the notion of PF, and proving that it achieves an approximation factor of at least 0.622.

In Section 4, we consider the objective of (approximately) maximizing the fraction of her PF utility that each bidder receives. Subsection 4.1 studies the set of instances involving just two items, for which we present a truthful mechanism which guarantees that all bidders will receive exactly the same fraction of their PF utility; this fraction is always at least $\frac{n}{n+1}$, where n is the number of bidders. For small values of n , we also describe more elaborate mechanisms achieving better approximation ratios: 0.828 instead of 2/3 for $n = 2$ and 0.775 instead of 3/4 for $n = 3$, both of which appear in Appendix B. Then, in Subsection 4.2, we address the general setting with arbitrary numbers of bidders and items. We present a truthful mechanism that performs increasingly well as the PF prices² increase. More specifically, if p_j^* is the PF price of item j , then the approximation factor guaranteed by this mechanism is equal to $\min_j (p_j^*/\lceil p_j^* \rceil)$. It is interesting to note that scenarios such as the privatization auction mentioned above involve a number of bidders much larger than the number of items; as a rule, we expect this to lead to high prices and a very good approximation of the participants’ PF utilities.

Related Work. Our setting is closely related to the large topic of fair division or cake-cutting [2, 16], which has been studied since the 1940’s, using the $[0, 1]$ interval as the standard representation of a cake. Each agent’s preferences take the form of a valuation function over this interval and the valuations of unions of subintervals are additive. The multiple divisible items setting that we focus on is equivalent to restricting the agents’ valuations to piecewise constant functions over the $[0, 1]$ interval. Some of the most common notions of fairness that have been studied in this literature are, proportionality, envy-freeness, and equitability [2, 16].³

²The prices induced by the market equilibrium when all bidders have a unit of budget.

³It is worth distinguishing the notion of PF from that of proportionality by noting that the latter is a much weaker notion, directly implied by the former. Also, note that all the mechanisms that we propose yield envy-free outcomes.

Despite the extensive work on fair resource allocation, truthfulness considerations have not played a major role in this literature. Most results related to truthfulness were weakened by the assumption that each agent would be truthful in reporting its valuations unless this strategy was dominated. Very recent work [5, 13, 18] studies truthful cake cutting variations using the standard notion of truthfulness according to which an agent may not be truthful unless doing so is a dominant strategy. Chen et al. [5] study truthful cake-cutting with agents having piecewise uniform valuations and they provide a polynomial-time mechanism that is truthful, proportional, and envy-free. They also design randomized mechanisms for more general families of valuation functions, while Mossel and Tamuz [13] prove the existence of truthful (in expectation) mechanisms satisfying proportionality in expectation for general valuations. Zivan et al. [18] aim to achieve envy-free Pareto optimal allocations of multiple divisible goods while reducing, but not eliminating, the agents' incentives to lie. The extent to which untruthfulness is reduced by their proposed mechanism is only evaluated empirically and depends critically on their assumption that the resource limitations are soft constraints. Our work is closely related to the recent papers of Guo and Conitzer [9] and of Han et al. [10], who also consider the truthful allocation of multiple divisible goods; their goal is to maximize the social welfare (We provide more details in Section 3).

Most of the papers mentioned above contribute to our understanding of the trade-offs between either truthfulness and fairness, or truthfulness and social welfare. Another direction that has been actively pursued is to understand and quantify the interplay between fairness and social welfare. Caragiannis et al. [4] measured the deterioration of the social welfare caused due to different fairness restrictions, the price of fairness. More recently, Cohler et al. [6] designed algorithms for computing allocations that (approximately) maximize social welfare while satisfying envy-freeness.

Our results fit into the general agenda of approximate mechanism design without money, explicitly initiated by Procaccia et al. [15]. The use of artificial currency aiming to achieve truthfulness and fairness in such scenarios, and the notion of *competitive equilibrium with equal incomes* was recently revisited by Budish [3] and Budish et al. [14].

Finally, the work of Zhang [17], despite the fact that it does not explicitly refer to the use of artificial currency, enforces that all the participants' budgets be spent, thus moving to an equivalent scenario. Unlike our work, they focus on evaluating the equilibrium points of the resource allocation game rather than on designing truthful mechanisms.

2 Preliminaries

Let M denote the set of m items and N the set of n bidders. Each bidder $i \in N$ has a valuation v_{ij} for each item $j \in M$ and each item is divisible, meaning that it can be cut into arbitrarily small pieces and then allocated to different bidders. The bidder valuations are scaled so that $\sum_j v_{ij} = 1$ for each bidder i .

Let x be an allocation of the resources among the bidders and $v_i(x)$ be the valuation of bidder i for this allocation. Such an allocation is *Proportionally Fair* (PF) if it is feasible, and for any other feasible allocation x' the aggregate proportional change to the valuations is not positive, i.e.:

$$\sum_{i \in N} \frac{v_i(x') - v_i(x)}{v_i(x)} \leq 0.$$

We focus on problems where every bidder has additive linear valuations; this means that, if a bidder is allocated a fraction f_j of each item j that she values at v_j , then her valuation for that

allocation will be $\sum_j f_j v_j$. For such additive linear valuations we can compute a PF allocation in polynomial time by assuming that each bidder has a *unit budget* of some artificial currency and then using the Eisenberg-Gale convex program to compute the market equilibrium of the induced market [7, 11]. This outcome need not be unique, but it provides unique item prices and achieved bidder valuations [8]; we will refer to these induced item prices as PF prices.

Given a valuation bid vector from each bidder (one bid for each item), we want to design mechanisms that output an allocation of items to bidders. We restrict ourselves to truthful mechanisms, i.e. mechanisms such that any false bid from a bidder will never return her a more valuable allocation. In designing such mechanisms we consider two objectives:

- The *SW objective*, which aims to output an allocation x (approximately) maximizing the social welfare, denoted $SW(x) = \sum_i v_i(x)$.
- The *PF objective*, which aims to output an allocation x (approximately) maximizing the value of $\rho(x) = \min_{i \in N} \left(\frac{v_i(x)}{v_i(x_{PF})} \right)$, where x_{PF} denotes the PF allocation.

Since maximizing these objectives via truthful mechanisms is infeasible in our setting, we will measure the performance of our mechanisms based on the extent to which they approximate them. More specifically, when referring to an approximation factor for the SW objective, this will be the minimum value of the ratio $SW(x)/SW(x^*)$ across all the relevant problem instances, where x is the output of the mechanism and x^* is the allocation that maximizes SW. For the PF objective the approximation factor will be the minimum value of $\rho(x)$ across all relevant instances.

3 Social Welfare Approximation

The problem of maximizing the SW objective in this model was first approached by Guo and Conitzer [9], who focused on the interesting case of two-bidder instances, which often draws the attention of researchers studying efficient or fair allocation. They first showed that no truthful mechanism can achieve better than a 0.841 approximation of the optimal social welfare, even for two items. They then studied a subclass of swap-dictatorial mechanisms and showed that no mechanism from that subclass can achieve better than a 0.5 approximation of the optimal social welfare when m , the number of items, is unbounded. Subsequent work of Han et al. [10] extended these negative results, showing that no swap-dictatorial mechanism can achieve better than a 0.5 approximation for the two-bidder case, and that for the general setting with $n > 2$, no truthful mechanism can achieve better than the trivial $1/m$ approximation. The main open question that remains in this setting is whether interesting truthful mechanisms for the two-bidder case exist beyond the class of swap-dictatorial mechanisms and whether such mechanisms can achieve an approximation factor better than 0.5.

We provide a positive answer to both open questions using the notion of PF as a tool. We first show that, for two-bidder instances, PF allocations approximate the SW objective very efficiently, even in the worst case. We then give two truthful mechanisms, one of which is inherently connected to the PF allocation, and show that combining the two gives a 0.622 approximation of the SW objective, beating the lower bound of 0.5 that was shown in [10] for all swap-dictatorial mechanisms.

3.1 The Social Welfare of Proportionally Fair Allocations

We start by showing that for two agents A, B , and multiple items, the social welfare of the PF allocation x_{PF} is a very good approximation of the social welfare achieved by x^* , the optimal allocation. Specifically:

Theorem 1. *For two agents and multiple items, $\frac{SW(x_{PF})}{SW(x^*)} \geq \frac{2\sqrt{3}+3}{4\sqrt{3}} \approx 0.933$.*

Proof. Assume that the items are ordered in a decreasing fashion w.r.t. the $\frac{v_{Aj}}{v_{Bj}}$ ratio. It is easy to verify that x^* allocates each item j either to agent A , or to agent B , depending on whether v_{Aj} or v_{Bj} is larger. We assume that items valued equally by both are allocated to agent B , so x^* can be thought of as defining an item e in the ordering mentioned above such that all the items preceding item e are allocated to agent A and all the rest to agent B . This type of “frontier” within this ordering, defining each player’s allocation, actually occurs with all Pareto efficient allocations, so the PF allocation also defines such a frontier⁴. W.l.o.g. we assume that, in the items’ ordering, the frontier of PF comes before the frontier of the social welfare maximizing allocation.

These two frontiers separate the set of items into three groups. Slightly abusing notation, let v_{Ag} and v_{Bg} denote the valuations of agents A and B respectively for group $g \in \{1, 2, 3\}$, where the group number g is consistent with the initial item ordering. Note that $v_{A1}/v_{B1} \geq v_{A2}/v_{B2}$, and $v_{A2}/v_{B2} \geq v_{A3}/v_{B3}$. The ratio that we are studying can thus be rewritten as follows:

$$\frac{SW(x_{PF})}{SW(x^*)} = \frac{v_{A1} + v_{B2} + v_{B3}}{v_{A1} + v_{A2} + v_{B3}} = 1 - \frac{v_{A2} - v_{B2}}{v_{A1} + v_{A2} + v_{B3}}. \quad (1)$$

Let $v_{A2} = kv_{B2}$ for some $k > 1$. Then, $v_{A1} \geq kv_{B1}$. Thus $k(v_{B1} + v_{B2}) \leq v_{A1} + v_{A2} \leq 1$, or $k(1 - v_{B3}) \leq 1$, which yields that $v_{B3} \geq (k-1)/k$. Also, by the definition of PF, since the PF solution allocates the second group of items to agent B , $\frac{v_{A2}}{v_{A1}} \leq \frac{v_{B2}}{v_{B2}+v_{B3}}$ which, after substituting for v_{A2} , yields $v_{A1} \geq k(v_{B2} + v_{B3})$. Since $v_{B2} + v_{B3} = 1 - v_{B1}$, this inequality can be rewritten as $v_{A1} \geq k(1 - v_{B1})$. Adding this inequality to $v_{A1} \geq kv_{B1}$ yields $v_{A1} \geq k/2$. Using these lower bounds for v_{A1} and v_{B3} in Equation (1), we get:

$$\frac{SW(x_{PF})}{SW(x^*)} \geq 1 - \frac{(k-1)v_{B2}}{\frac{k}{2} + kv_{B2} + \frac{k-1}{k}} = 1 - \frac{2k(k-1)v_{B2}}{2k^2v_{B2} + k^2 + 2k - 2}.$$

We can then show that for any value of v_{B2} , the right hand side of the inequality is minimized for k as large as possible, and that k is restricted to be at most $\frac{1}{v_{B2}+0.5}$, implying that $k = \frac{1+\sqrt{3}}{2}$ in the worst case, thereby proving the theorem (details in Appendix A). \square

3.2 Two Truthful Mechanisms

Swap-Dictatorial Mechanism. Our first mechanism picks one of the two players with probability $\frac{1}{2}$, and then, based on her bid, provides her with her most preferred bundle of at most $m/2$ items. The other player is allocated everything that is left. Clearly, this mechanism is truthful since a bidder would prefer to bid truthfully if picked, and her bid does not affect the allocation if she is not picked. This mechanism achieves a 0.5 approximation: If v is some bidder’s valuation

⁴If the PF solution dictates that an item has to be shared between the two agents, we can just split it into two items such that each one is allocated to a different agent; such a change does not affect the item ordering.

for her allocated bundle if she is picked by the mechanism, then the least value that she receives if she is not picked (and the other bidder is allocated her favorite bundle) is $1 - v$. Therefore, both bidders get an expected value of at least 0.5, yielding an expected social welfare of at least 1; the optimal social welfare is of course at most 2. This is tight, as the following example shows: there are 4 items valued by bidder A at $(1 - 2\epsilon, \epsilon, \epsilon/2, \epsilon/2)$ and by Bidder B at $(\epsilon, 1 - 2\epsilon, \epsilon/2, \epsilon/2)$, for some small $\epsilon > 0$; the mechanism achieves a $1/2 + \Theta(\epsilon)$ approximation. Finally, we note that we can use a deterministic version of this mechanism: it plays this game twice, once with Bidder A as the dictator and once with Bidder B in this role. Each game is played with a set comprising one half of each original item.

Partial Allocation Mechanism. We now present an interesting non-swap-dictatorial truthful mechanism which we call the *Partial Allocation* (PA) mechanism. Let $v_A = v_A(x_{PF})$ and $v_B = v_B(x_{PF})$ denote respectively the fraction of their total valuation that Bidders A and B receive in the PF allocation (remember that the total valuation is equal to 1). The mechanism allocates Bidder A a fraction v_B of each of the items in her PF allocation and, similarly, Bidder B a fraction v_A of each of the items in her PF allocation.⁵ Note that, as a direct implication, Bidder A's utility is a v_B approximation of her PF utility and similarly Bidder B's utility is a v_A approximation. In contrast to the previous mechanism, the types of instances for which this mechanism performs poorly are the ones where, for example, both bidders value all items equally; this would cause both bidders to receive only a 0.5 fraction of their PF allocation, a 0.5 approximation of the SW objective.

Lemma 1. *The PA mechanism is truthful.*

Proof. We show that Bidder A is truthful; the same argument applies to Bidder B. We rescale the valuations of Bidder A so that her new valuation for the PF allocation is $\bar{v}_A = v_B$. Suppose Bidder A changes her bid and, in the changed PF allocation, gains allocation of value g_A and loses value l_A , according to the scaled valuations. Let the gains and losses to Bidder B have values g_B and l_B . Since $\bar{v}_A = v_B$ and given the definition of PF, we get $l_B \geq g_A$ and $l_A \geq g_B$, for otherwise either one of these shifts of allocation between A and B would have lead to a higher percentage increase than decrease. The rescaled valuation of Bidder A after changing her bid would be:

$$\begin{aligned}
(v_B - l_B + g_B)(\bar{v}_A - l_A + g_A) &= v_B\bar{v}_A + v_B(g_A - l_A) + \bar{v}_A(g_B - l_B) + (g_B - l_B)(g_A - l_A) \\
&= v_B[\bar{v}_A - (l_A - g_A) - (l_B - g_B)] - (g_A - l_A)(l_B - g_B) \\
&= v_B[\bar{v}_A - (l_B - g_A) - (l_A - g_B)] - (g_A - l_A)(l_B - g_B) \\
&\leq v_B\bar{v}_A - (g_A - l_A)(l_B - g_B).
\end{aligned}$$

Now, if $g_A \geq l_A$, then as $l_B \geq g_A$, and $l_A \geq g_B$, it follows that $l_B \geq g_B$; thus the new valuation is at most the original $v_B \cdot \bar{v}_A$. A similar argument applies if $g_B \geq l_B$. Otherwise, $g_A \leq l_A$ and $g_B \leq l_B$ and the reduction in value for Bidder A as is at least $v_B[(l_A - g_A) + (l_B - g_B)] + (g_A - l_A)(l_B - g_B) = (l_A - g_A)[v_B - (l_B - g_B)] + (l_B - g_B) \geq 0$, as $l_B \leq v_B$. \square

3.3 The Hybrid Mechanism

We now provide a non-swap-dictatorial truthful mechanism that outperforms all swap-dictatorial mechanisms, giving a 0.622 approximation of the SW objective. Our mechanism is a combination

⁵Like before, we can split an item into two in order to avoid any sharing of items in the PF allocation.

of the swap-dictatorial mechanism and the PA mechanism described in the previous subsection. A randomized version of it would run the swap-dictatorial mechanism with probability 0.5 and the PA mechanism with the remaining 0.5 probability. Again, this can be implemented in a deterministic fashion by splitting all items in half and using a different mechanism for each set of halves. The benefit of combining them comes from the fact that one performs well when the other one does not.

Lemma 2. *The hybrid mechanism outputs an allocation x_m always satisfying $\frac{SW(x_m)}{SW(x_{PF})} \geq \frac{2}{3}$.*

Proof. (Sketch.) The swap-dictatorial mechanism always provides a social welfare of at least 1. The PA mechanism provides a social welfare of $2v_A v_B$. Thus, $\frac{SW(x_m)}{SW(x_{PF})} = \frac{0.5 + v_A v_B}{v_A + v_B}$. We show that for all $(v_A, v_B) \in [0.5, 1]^2$ this is minimized for $v_A = 0.5$ and $v_B = 1$ (details in Appendix A). \square

This lemma combined with Theorem 1 immediately implies:

Theorem 2. *The hybrid mechanism always satisfies $\frac{SW(x_m)}{SW(x^*)} \geq 0.933 \cdot \frac{2}{3} \approx 0.622$.*

4 Proportional Fairness Approximation

Despite the fact that maximizing social welfare is a very natural objective, one can quickly verify that allocations that are very efficient in terms of social welfare can be very unfair to some bidders. Dealing with problems such as the one that arose with the Czech privatization auctions [1] calls for solutions that are fair. In what follows, we use the PF allocation as a benchmark for the “fair share” that each bidder should be receiving and we aim to provide every bidder with a good approximation of his value for that share. We start by focusing on the case of two items which helps build some intuition for our solution to the general case that follows.

4.1 Two Items

Proportionally fair allocation for two items. For instances involving just two items t and b (or top and bottom), for simplicity we choose one of the two items (w.l.o.g. b) and rescale the valuations of each bidder so that they all have a valuation of 1 for item b . (If a bidder has zero valuation for item b , then her valuation for item t is set to be ∞ .) We then sort the bidders in decreasing order of their valuation for item t (breaking ties arbitrarily). The proof of the following lemma (in Appendix A) shows that for two-item instances there always exists a PF allocation such that at most one bidder is allocated parts of both items. The proof also shows that if such a bidder exists and her valuation for item t is v , then her valuation for the PF allocation will be $\frac{v+1}{n}$. Such a bidder will also be defining the relative value of the PF prices of the two items (the PF price of t is v times that of b), so we will henceforth call her the *Ratio Defining Bidder*, denoted R_b .

Lemma 3. *For two item instances, there always exists a PF allocation with at most one R_b .*

Mechanism for many bidders and two items. In this scheme *every* bidder receives a fraction of just one item. The mechanism can be thought of as a tie breaking rule that doesn’t allow the existence of a bidder being allocated parts of more than one item. More specifically, given the bidders’ bids, the mechanism computes a PF allocation with at most one R_b ⁶ and if one exists for

⁶This can be done by using binary search over the bidder ordering until R_b , if one exists, is found.

that allocation, we instead force her to either equally share item t with the TOP bidders or item b with the BOTTOM bidders (respectively the bidders before and after R_b in the above ordering). Among these two options, the mechanism chooses the one that maximizes R_b 's utility. Let ρ denote the fraction of her PF utility that R_b receives. Every other bidder's allocation is reduced as needed so as to provide that same ρ fraction. We call this the *Single Item* (SI) mechanism.

Lemma 4. *With n bidders, the SI mechanism achieves an approximation factor of $\frac{n}{n+1}$.*

Proof. The worst ρ occurs when R_b achieves the same utility with either option. If R_b is the k -th bidder in the ordering, this means that she values being allocated $1/k$ of the top item equally with $\frac{1}{n-k+1}$ of the bottom item, i.e. $\frac{v}{k} = \frac{1}{n-k+1}$, or $v = \frac{k}{n-k+1}$. The approximation factor then, on substituting for v , is given by:

$$\frac{v/k}{(v+1)/n} = \frac{1/(n-k+1)}{[(n+1)/(n-k+1)]/n} = \frac{n}{n+1}.$$

□

Lemma 5. *The SI mechanism is truthful. (Proof in Appendix A.)*

This mechanism performs really well as the number of bidders increases, but can be further improved for small values of n . In Appendix B we present an interesting and more efficient mechanism for the special case of $n = 2$, which achieves an approximation factor of 0.828 instead of $2/3$, and a similar mechanism for $n = 3$, achieving an approximation factor of 0.775 instead of $3/4$. As n increases beyond $n = 3$, such more elaborate schemes provide, at best, very modest gains compared to the $\frac{n}{n+1}$ approximation mechanism presented above.

4.2 Many Bidders and Many Items

Using the intuition acquired from the two-item case, we now describe the *Strong Demand Matching* mechanism (SDM) for the general case. Informally speaking, SDM provides every bidder with a unit budget and then aims to discover item prices such that the demand of each bidder can be satisfied using (a fraction of) just one item. This approach resembles the SI mechanism, but the structure of the PF solution can now be much more complicated. This mechanism tackles the problem of allocating a collection of goods for the very natural set of problem instances for which the bidders may have arbitrary valuation functions, yet no item is undemanded if its price is low. In what follows we describe the process by which SDM increases the prices of overdemanded items in a fashion that maintains truthfulness and yields prices that are very close to the PF prices.

Let p_j denote the price of item e_j , and let the *bang per buck* that Bidder i gets from item e_j equal v_{ij}/p_j . We say that item e_j is an MBB item of Bidder i if Bidder i gets the maximum bang per buck from that item⁷. For a given price vector p , let the demand graph $D(p)$ be a bipartite graph with bidders on one side and items on the other, such that an edge between Bidder i and item e_j exists if and only if e_j is an MBB item of Bidder i . We call $c_j = \lfloor p_j \rfloor$ the *capacity* of item e_j when its price is p_j , and we say an assignment of bidders to items is *valid* if it matches each bidder to at most one item and no item e_j is matched to more than c_j bidders. Given a valid assignment A , we say an item e_j is *reachable* from Bidder i if there exists an alternating path $(i, j_1, i_1, j_2, i_2, \dots, j_k, i_k, j)$ in the graph $D(p)$ such that edges $(i_1, j_1), \dots, (i_k, j_k)$ lie in the

⁷Note that for each bidder there could be multiple MBB items.

assignment A . Finally, let $d(R)$ be the collection of bidders with all their MBB items in set R .

The SDM mechanism initializes all item prices to $p_j = 1$ and iterates as follows:

1. Find a valid assignment that maximizes the number of matched bidders.
If all the bidders are matched, conclude with Step 3.
2. Let U be the set of bidders who are not matched in Step 1.
Let R be the set of all items reachable from bidders in the set U .
Raise the price of each item e_j in R from p_j to $x \cdot p_j$,
where $x \geq 1$ is the minimum value for which one of the following events takes place:
 - (a) The price of an item in R reaches an integral value. If this happens, repeat Step 1.
 - (b) For some bidder $b_i \in d(R)$, her set of MBB items increases, causing R to grow:
 - i. If for each item e_j added to R , the number of bidders matched to it equals c_j ,
continue with Step 2.
 - ii. If some item e_j added to R has c_j greater than the number of bidders matched to it,
continue with Step 1.
3. Every bidder matched to some item e_j is allocated a fraction $1/p_j$ of that item.

It remains to explain how to carry out Step 2. Set R can be found using a breadth-first-search like algorithm. To determine when (a) is reached, we just need to know the smallest $\lceil p_j \rceil / p_j$ ratio over all items whose price is being increased. For (b), we need to calculate, for each bidder in $d(R)$, the ratio of the *bang per buck* for her MBB items and for the items outside the set R .

Running time. If $c(R) = \sum_{j \in R} c_j$ denotes the total capacity in R , it is not difficult to see that if U is non-empty, $|d(R)| > c(R)$. Note that each time either event (a) or event (b)-ii occurs, $c(R)$ increases by at least 1, and thus, using the alternating path from a bidder in the set U to the corresponding item, we can increase the number of matched bidders by at least 1; this means that this can occur at most n times. The only other events are the unions resulting from (b)-i. There can be at most $\min(n, m)$ of these, and they are followed by either Step (a) or (b)-ii. Thus there are $O(n * \min(n, m))$ iterations of Step (b)-i and $O(n)$ iterations of Steps 1 and (b)-ii.

Correctness. Let p^* be the PF prices and let q be the prices computed by the algorithm.

Lemma 6. *Let $f = \max_j \lceil p_j^* \rceil / p_j^*$. Then $q \leq fp^*$.*

Proof. First note that at prices fp^* , the MBB items for each bidder are the same as at prices p^* . It is not difficult to see that at prices fp^* every bidder can be allocated to exactly one item from among her MBB items such that the number of bidders allocated to an item e_j is less than or equal to fp_j^* . To show this, consider a PF allocation. Form the following graph on items and bidders — add an edge between a bidder and an item if a portion of this item is assigned to this bidder in the PF solution. If there exists a cycle in this graph, one can remove an edge in this cycle by reallocating along the cycle while maintaining the utility of every bidder. Hence there is a PF allocation in which this graph is a forest. Now for a given tree, root it at an arbitrary bidder. For

each Bidder b in this tree, assign it to one of its child items, if any, and otherwise to its parent. The result is that for each item e_j , at most $\lceil p_j \rceil$ bidders will be assigned to it.

Now, suppose that some $q_j > fp_j^*$. Consider the first time t at which some price q_i starts to increase from fp_i^* . Let S be the set of items e_i whose price is currently fp_i^* . We will show that no item in set S will be part of the set R , and hence the prices of these items will not increase. Let T_q and T_{fp^*} be the sets of bidders who have edges to some item in the set S at the current prices and at prices fp^* , respectively. Clearly $T_q \subseteq T_{fp^*}$. Also suppose a bidder $b \in T_{fp^*}$ has an edge to an item outside set S at prices fp^* ; this means that $b \notin T_q$ at the current prices as b will strictly prefer the item outside the set S . Thus if $b \in T_q$, this implies that b has no edges outside the set S at prices fp^* , and so b was allocated to some item in set S at prices fp^* . Since we know that at prices fp^* all the bidders can be allocated to some item, this implies that $|T_q| \leq c(S)$. Thus even at current prices, all the bidders in T_q can be allocated to items in set S , and hence no item in set S can be part of the set R . \square

Truthfulness. We argue by contradiction. Suppose that some Bidder b were not truthful in the above algorithm, which we name algorithm \mathcal{A} . First, we consider an alternate algorithm \mathcal{A}' , and show that it is a dominant strategy for b to be truthful in \mathcal{A}' . Then we show that \mathcal{A} and \mathcal{A}' produce the same outcomes, and consequently b should also be truthful in \mathcal{A} . \mathcal{A}' proceeds as follows. It begins by running algorithm \mathcal{A} but with b absent (the first run), yielding prices p' . Then it runs \mathcal{A} on all n bidders, but starting from prices p' (the second run).

Lemma 7. *b is truthful in algorithm \mathcal{A}' .*

Proof. Suppose the second run of the algorithm ends when b can be matched using an alternating path that ends at an item e . Suppose at this point, the price of an item $e_j \in R$ is p_j . It is easy to see that bidder b has no incentives to lie to obtain an item that is not in R as the prices of these items is completely defined by other bidders. Suppose that by lying, Bidder b is able to get an item $e_j \in R$ at a price $p'_j < p_j$. Suppose that this happens with an alternating path that starts at b and ends at some item e' . Now if this path existed in the truthful scenario when the price of item e_j reaches p'_j , then even in the truthful scenario b would have been matched when the price of item e_j is p'_j . Thus this path doesn't exist in the truthful scenario when the price of item e_j is p'_j . But why does this path not exist in the truthful scenario? It must be the case that some item on this path has a higher price than the price in the lying scenario. A higher price on this item means that in the truthful scenario item e_j would have been matched via some alternating path ending at e' before the price of item e_j reached p'_j , a contradiction. \square

Lemma 8. *\mathcal{A} and \mathcal{A}' have the same outcome.*

The proof of this Lemma is similar to that of Lemma 6 and is deferred to Appendix A.

Corollary 1. *\mathcal{A} is truthful.*

Theorem 3. *The SDM mechanism achieves an approximation factor of $\rho = \min_j (p_j^*/\lceil p_j^* \rceil)$. If $\min_j p_j^* = k$, this is an approximation factor of at least $\frac{k}{k+1}$.*

Proof. If a bidder is allocated a portion of item e_j , she receives a $1/q_j$ fraction. But by Lemma 6, $1/q_j \geq 1/(fp_j^*) \geq 1/\lceil p_j^* \rceil$. In value, her PF allocation equals a $1/p_j^*$ fraction of item e_j . Thus she achieves an approximation factor of $p_j^*/\lceil p_j^* \rceil$. The result follows on minimizing over all bidders. \square

References

- [1] R. Aggarwal and J.T. Harper, *Equity valuation in the czech voucher privatization auctions*, Financial Management (Winter 2000), 77–100.
- [2] S. Brams and A. Taylor, *Fair division: from cake cutting to dispute resolution*, Cambridge University Press, Cambridge, 1996.
- [3] E. Budish, *The combinatorial assignment problem: approximate competitive equilibrium from equal incomes*, BQGT '10, ACM, 2010, pp. 74:1–74:1.
- [4] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou, *The efficiency of fair division*, Theory Comput. Syst. **50** (2012), no. 4, 589–610.
- [5] Y. Chen, J.K. Lai, D.C. Parkes, and A.D. Procaccia, *Truth, justice, and cake cutting*, AAAI, 2010.
- [6] Y.J. Cohler, J.K. Lai, D.C. Parkes, and A.D. Procaccia, *Optimal envy-free cake cutting*, AAAI, 2011.
- [7] N.R. Devanur, C.H. Papadimitriou, A. Saberi, and V.V. Vazirani, *Market equilibrium via a primal-dual-type algorithm*, FOCS, 2002, pp. 389–395.
- [8] E. Eisenberg and D. Gale, *Consensus of subjective probabilities: The pari-mutuel method*, Ann. Math. Stat. **30** (1959), 165–168.
- [9] M. Guo and V. Conitzer, *Strategy-proof allocation of multiple items between two agents without payments or priors*, AAMAS, 2010, pp. 881–888.
- [10] L. Han, C. Su, L. Tang, and H. Zhang, *On strategy-proof allocation without payments or priors*, WINE, 2011, pp. 182–193.
- [11] K. Jain and V.V. Vazirani, *Eisenberg-gale markets: algorithms and structural properties*, STOC, 2007, pp. 364–373.
- [12] F. Kelly, *Proportional fairness*, www.statslab.cam.ac.uk/~frank/pf/.
- [13] E. Mossel and O. Tamuz, *Truthful fair division*, SAGT, 2010, pp. 288–299.
- [14] A. Othman, T. Sandholm, and E. Budish, *Finding approximate competitive equilibria: efficient and fair course allocation*, AAMAS '10, 2010, pp. 873–880.
- [15] A.D. Procaccia and M. Tennenholtz, *Approximate mechanism design without money*, ACM Conference on Electronic Commerce, 2009, pp. 177–186.
- [16] J.M. Robertson and W.A. Webb, *Cake-cutting algorithms - be fair if you can*, A K Peters, 1998.
- [17] L. Zhang, *The efficiency and fairness of a fixed budget resource allocation game*, ICALP, 2005, pp. 485–496.
- [18] R. Zivan, M. Dudík, S. Okamoto, and K.P. Sycara, *Reducing untruthful manipulation in envy-free pareto optimal resource allocation*, IAT, 2010, pp. 391–398.

A Omitted Proofs

of Theorem 1. Assume that the items are ordered in a decreasing fashion w.r.t. the $\frac{v_{Aj}}{v_{Bj}}$ ratio. It is easy to verify that x^* allocates each item j either to agent A , or to agent B , depending on whether v_{Aj} or v_{Bj} is larger. We assume that items valued equally by both are allocated to agent B , so x^* can be thought of as defining an item in the ordering mentioned above such that all the items preceding this item are allocated to agent A and all the rest to agent B . This type of “frontier” within this ordering, defining each player’s allocation, actually occurs with all Pareto efficient allocations, so the PF allocation also defines such a frontier⁸. W.l.o.g. we assume that the frontier of PF comes before the frontier of the social welfare maximizing allocation in the items’ ordering.

These two frontiers separate the set of items into three groups. Slightly abusing notation, let v_{Ag} and v_{Bg} denote the valuations of agents A and B respectively for group $g \in \{1, 2, 3\}$, where the group number g is consistent with the initial item ordering. Note that $v_{A1}/v_{B1} \geq v_{A2}/v_{B2}$, and $v_{A2}/v_{B2} \geq v_{A3}/v_{B3}$. The ratio that we are studying can thus be rewritten as follows:

$$\frac{SW(x_{PF})}{SW(x^*)} = \frac{v_{A1} + v_{B2} + v_{B3}}{v_{A1} + v_{A2} + v_{B3}} = 1 - \frac{v_{A2} - v_{B2}}{v_{A1} + v_{A2} + v_{B3}}. \quad (2)$$

Let $v_{A2} = kv_{B2}$ for some $k > 1$. Then, $v_{A1} \geq kv_{B1}$. Also, by the definition of PF, since the PF solution allocates the second group of items to agent B , then $\frac{v_{A2}}{v_{A1}} \leq \frac{v_{B2}}{v_{B2} + v_{B3}}$ which, after replacing for v_{A2} , gives $v_{A1} \geq k(v_{B2} + v_{B3})$. Since $v_{B2} + v_{B3} = 1 - v_{B1}$, this inequality can be rewritten as $v_{A1} \geq k(1 - v_{B1})$. Adding this inequality to $v_{A1} \geq kv_{B1}$ yields $v_{A1} \geq k/2$. Also, according to the definition of k , we get that $k(v_{B1} + v_{B2}) \leq v_{A1} + v_{A2} \leq 1$, or $k(1 - v_{B3}) \leq 1$, which yields that $v_{B3} \geq (k - 1)/k$. Using these lower bounds for v_{A1} and v_{B3} in Equation (2), we get:

$$\frac{SW(x_{PF})}{SW(x^*)} \geq 1 - \frac{(k - 1)v_{B2}}{\frac{k}{2} + kv_{B2} + \frac{k-1}{k}} = 1 - \frac{2k(k - 1)v_{B2}}{2k^2v_{B2} + k^2 + 2k - 2}. \quad (3)$$

The lower bound implied by this inequality is minimized when the fraction on the right hand side is maximized. Assuming that v_{B2} is fixed, we take the partial derivative w.r.t. k , which is equal to:

$$\left(\frac{2k(k - 1)v_{B2}}{2k^2v_{B2} + k^2 + 2k - 2} \right)'_k = \frac{(2k^2v_{B2} + 3k^2 - 4k + 2)2v_{B2}}{(2k^2v_{B2} + k^2 + 2k - 2)^2}.$$

It is easy to verify that this is positive because $3k^2 - 4k + 2 > 0$ for any value of k . This means that for any value of v_{B2} , the fraction is maximized when k is as large as possible. But we know that $kv_{B2} = v_{A2} \leq 1 - v_{A1}$, and since $v_{A1} \geq k/2$, this yields $k \leq \frac{1}{v_{B2} + 0.5}$. Thus to maximize the fraction we let $k = \frac{1}{v_{B2} + 0.5}$, or $v_{B2} = \frac{2-k}{2k}$. Substituting for v_{B2} in Inequality (3) yields:

$$\frac{SW(x_{PF})}{SW(x^*)} \geq 1 - \frac{-k^2 + 3k - 2}{4k - 2},$$

with the right hand side minimized for $k = \frac{1+\sqrt{3}}{2}$, which proves the theorem. \square

⁸If the PF solution dictates that an item has to be shared between the two agents, we can just split it into two items such that each one is allocated to a different agent without affecting the ordering.

of Lemma 2. Note that the dictatorial mechanism always provides a social welfare of at least 1 if used (each agent gets a value of exactly 0.5). The second mechanism provides a social welfare of $2v_A v_B$. We can therefore express the ratio of the social welfare of an allocation x_m that is the outcome of combining these two mechanisms, over the social welfare of the PF allocation as $\frac{SW(x_m)}{SW(x_{PF})} = \frac{0.5+v_A v_B}{v_A+v_B}$.

We only need to show that $\frac{0.5+v_A v_B}{v_A+v_B} \geq \frac{2}{3}$ for all v_A, v_B combinations in $[0.5, 1]^2$ (since it must be the case that $v_A, v_B \geq 0.5$ in the PF allocation). We treat v_B as a given constant and take the derivative w.r.t. v_A . This gives a derivative of $\frac{v_B^2 - 0.5}{(v_A+v_B)^2}$, which is negative for $v_B^2 < 0.5$, positive for $v_B^2 > 0.5$, and zero for $v_B^2 = 0.5$. This means that, given a value of $v_B < \sqrt{0.5} \approx 0.7$, the ratio is minimized for $v_A = 1$, and given a value of $v_B > \sqrt{0.5} \approx 0.7$, the ratio is minimized for $v_A = 0.5$. Given the symmetry of the ratio w.r.t. to v_A and v_B , the exact same argument applies for the values of v_B minimizing the ratio given a value of v_A . We conclude that the ratio is minimized either when $v_A = 0.5$ and $v_B = 1$ or when $v_A = 1$ and $v_B = 0.5$. In either case, the ratio evaluates to $\frac{2}{3}$, proving the theorem. \square

of Lemma 3. Let the PF prices be p_t and p_b , and note that, since $p_t + p_b = n$, either both prices are integers or neither is. If both prices are integers, it is easy to see that the following allocation is PF: the first p_t bidders are assigned a $1/p_t$ fraction of item t and the remaining p_b bidders are assigned a $1/p_b$ fraction of item b . If neither of the prices is an integer, then we get a PF allocation by giving each of the first $\lfloor p_t \rfloor$ bidders (the TOP bidders) a $1/p_t$ fraction of item t , Bidder $\lfloor p_t \rfloor$ a $\frac{p_t - \lfloor p_t \rfloor}{p_t}$ fraction of item t and a $\frac{p_b - \lfloor p_b \rfloor}{p_b}$ fraction of item b , and each of the remaining bidders (the BOTTOM bidders) a fraction of $1/p_b$ of item b .

Then, if a bidder being allocated a fraction of both the top and the bottom items in the PF solution exists and her valuation for the top item is v , notice that in that PF solution the allocation of every TOP bidder comprises a fraction of the top item alone; similarly, the allocations for BOTTOM bidders come from the bottom item alone. If there are $k-1$ TOP bidders, then the price of item t in the PF solution would be equal to $k-1+x$, and that of item b would be equal to $n-k+1-x$, where $x < 1$ is the amount Bidder k spends on item t . Notice that since Bidder k is interested in both items at these prices, this means that her valuations are the ones defining the ratio between the two prices (henceforth, we call her the *Ratio Defining Bidder*, denoted R_b), i.e.:

$$\frac{k-1+x}{n-k+1-x} = v \quad (4)$$

Thus:

$$x = \frac{(n-k+1)v - (k-1)}{v+1}. \quad (5)$$

Every TOP bidder gets a fraction of item t equal to the ratio of her budget over the final price of that item, i.e. a fraction equal to $\frac{v+1}{vn}$. Similarly, every BOTTOM bidder gets a fraction of the bottom item equal to $\frac{v+1}{n}$. Finally, R_b gets a fraction of item t equal to $\frac{n-k+1}{n} - \frac{(k-1)}{vn}$ and a fraction of item b equal to $\frac{k}{n} - \frac{(n-k)v}{n}$. Notice that if R_b were truthful, then this allocation would offer her utility equal to $\frac{v+1}{n}$, equal to that of every BOTTOM bidder and of every TOP bidder. This utility is also equal to the utility she would obtain from a $\frac{v+1}{vn}$ fraction of the top item, or a $\frac{v+1}{n}$ fraction of the bottom item. \square

of Lemma 5. To verify that this mechanism is truthful, we consider the strategies of any TOP bidder, any BOTTOM bidder and those of a possible ratio defining bidder.

TOP bidders (the same arguments apply for BOTTOM bidders as well):

- A TOP bidder overbids: then her allocation is unchanged.
- A TOP bidder underbids but continues to overbid R_b : then her allocation is unchanged.
- A TOP bidder underbids R_b and R_b had been forced to share the *top* item: then the only possible change is that the TOP bidder ends up sharing the bottom item which is something not even R_b wanted.
- A TOP bidder underbids R_b and R_b had been forced to share the *bottom* item. If R_b does not move then the ρ for the BOTTOM bidders is reduced, and hence it is reduced for everyone; while if R_b moves, then the best option for the TOP Bidder is to become the ratio defining bidder, and the only choice with the underbid is for the TOP Bidder to be sharing the bottom item. But her ρ ($= n/[(v+1)(n-k+1)]$) is smaller than the approximation of R_b which is the approximation she achieved before.

Ratio Defining Bidder (w.l.o.g. assume that she had been forced to share the top item):

- There is no point in overbidding.
- By underbidding, one possible outcome is that this bidder shares the bottom item along with those who had already been sharing it, which of course by definition is worse for her.
- By underbidding one of the BOTTOM bidders this bidder ceases to be ratio defining. The interesting case is when there is a new ratio defining bidder and she is forced to go up. In that case, the ρ for that bidder ($vn/k(v+1)$) is smaller than the approximation of the original ratio defining bidder and everyone suffers this same ρ .

□

of Lemma 8. Suppose not for a contradiction. W.l.o.g. suppose that some price is higher in the outcome of \mathcal{A}' (if not, switch the roles of \mathcal{A} and \mathcal{A}').

In the run of the mechanism \mathcal{A}' , consider a time when no item has exceeded the final price given by the mechanism \mathcal{A} but some items have reached that price. Let S be the set of items at this point that have prices equal to their final prices in \mathcal{A} . We will show that no item in set S will be part of the set R from this time onwards in the mechanism \mathcal{A}' , and hence the prices of these items will not increase in \mathcal{A}' . Let T' and T be the set of bidders who have edges to some item in the set S at the current prices in \mathcal{A}' , and at the final prices in \mathcal{A} , respectively. Clearly, $T' \subseteq T$. Also suppose a bidder $b \in T$ has an edge to an item outside set S at the final prices of \mathcal{A} ; this means that $b \notin T'$ at the current prices as b will strictly prefer the item outside the set S . Thus if $b \in T'$, this implies that b has no edges to an item outside the set S at the final prices given by \mathcal{A} , and so in \mathcal{A} , b was allocated to some item in set S . Since we know that at the final prices given by \mathcal{A} , all the bidders can be allocated to some item, this implies that $|T'| \leq c(S)$. Thus, even at current prices, all the bidders in T' can be allocated to items in the set S , and hence no item in set S can be part of the set R . Thus no item can have higher final price in \mathcal{A}' than in \mathcal{A} . A contradiction. □

B Improved Mechanisms for Two Items

B.1 Two Bidders

We start by assuming that both bidders are interested in both items, i.e. have non-zero valuation for both of them; we touch on the other cases at the end of this section.

Just as we did for the many bidder case, we scale their valuations so that both bidders assign a value of 1 to item b . Then, let u and v be the respective valuations of Bidder A and Bidder B for item t . W.l.o.g. we can assume that $u > 1$ and $u > v$, which means that Bidder A will never be R_b and Bidder B will be R_b only if $v > 1$. If no R_b exists, then the PF allocation is clearly truthful, but when Bidder B is the Ratio Defining Bidder, our 2-bidder 2-item mechanism defines the final allocation as a function of just v . More specifically, Bidder B gets a fraction $b(v) = \frac{1}{v}$ of item b and a fraction $t(v) = \frac{1}{2} - \frac{1}{2v^2}$ of item t . Finally, Bidder A is allocated all of item t that is not allocated to Bidder B. The intuition behind this mechanism is that, if Bidder B overbids regarding her valuation for item t , then she loses part of item b .

Theorem 4. *The 2-bidder, 2-item mechanism is truthful and achieves an approximation factor of $2 \cdot (\sqrt{2} - 1) \approx 0.828427$ of the PF objective.*

Proof. Notice that the valuation of Bidder B would equal $t(\bar{v})v + b(\bar{v})$ if her bid is \bar{v} , and for the mechanism to be truthful we need $t(\bar{v})v + b(\bar{v}) \leq t(v)v + b(v)$ for all \bar{v} . It is easy to verify that this holds for our mechanism since the partial derivative of the left hand side w.r.t. \bar{v} is equal to $\frac{v-\bar{v}}{\bar{v}^3}$, and hence v is the optimal solution for Bidder B.

Regarding the approximation ratio, the PF allocation gives a fraction of $\frac{1}{2} - \frac{1}{2v}$ of item t and all of item b to Bidder B, yielding an approximation factor of $\frac{v^2+1}{v^2+v}$, which is minimized for $v = 1 + \sqrt{2}$, giving $\rho = 2 \cdot (\sqrt{2} - 1) \approx 0.828427$. \square

Other cases. If just one bidder is interested in both items, we can view it as the previous case with $u = \infty$ for the bidder interested in only one item. If the bidders are each interested in only one item, if these are distinct items, the bidders are each allocated the item they want in full. If it is the same item, they each receive half of it. In all these cases, $\rho \geq 2 \cdot (\sqrt{2} - 1)$.

B.2 Three Bidders

Theorem 5. *There is a truthful mechanism for 3 bidders and 2 items achieving an approximation factor of $(12 - \sqrt{12})/11 \approx 0.77599$.*

Proof. Our goal is to find two functions $t(v)$ and $b(v)$ defining the fractions of the top and bottom items respectively that R_b will be assigned as a function of her valuation v for the top item. Before we move on to define such functions, we note that if they are functions of v alone, and if R_b obtains a fraction ρ of her utility in the PF solution, then every other bidder will have to obtain exactly the same fraction ρ of her PF utility. Were this not the case, we could easily construct an example where a bidder that gets a different approximation has valuation $v - \epsilon$ or $v + \epsilon$ (i.e. she is practically the same bidder as R_b). In that case, the bidder with the worse approximation would have the option of changing her bid by at most 2ϵ and securing the better approximation factor. Therefore every bidder must be offered the same fraction ρ of her PF allocation utility.

Now we consider the case when the middle bidder is R_b . We assume w.l.o.g. that this bidder prefers the top item ($v \geq 1$). Then, by Equation (4), $v \in [1, 2]$. Following the approach used for the case of two bidders, we consider the following families of functions:

$$t(v) = \alpha - \beta \frac{1}{v^2} \quad \text{and} \quad b(v) = \gamma \frac{1}{v} - \delta,$$

where α, β, γ and δ are non-negative constants which we will choose so as to maximize ρ while ensuring both truthfulness and that the resulting solution can be allocated. More specifically, for R_b to be truthful, the utility function $t(\bar{v})v + b(\bar{v})$ of R_b has to be maximized at $\bar{v} = v$, i.e. it is a necessary condition that:

$$t'(v)v + b'(v) = 0 \Rightarrow \beta \frac{2}{v^3}v - \gamma \frac{1}{v^2} = 0 \Rightarrow 2\beta = \gamma. \quad (6)$$

Using the fact that $2\beta = \gamma$, we can now reduce the number of constants to three by replacing γ . Then, ρ , the approximation factor, will, by definition, be equal to:

$$\rho = \frac{t(v)v + b(v)}{\frac{1}{3}(v+1)} \Rightarrow \rho = \frac{3(\alpha v^2 - \delta v + \beta)}{v^2 + v}. \quad (7)$$

Given $t(v)$ and $b(v)$, the fraction of the top item that remains for the TOP bidder is $1 - t(v)$ and the fraction of the bottom item that remains for the BOTTOM bidder is $1 - b(v)$. We need to make sure that these remaining fractions suffice to provide both the TOP and the BOTTOM bidders with a ρ -approximation of their PF utilities. These restrictions translate to:

$$\rho \frac{v+1}{3v} \leq 1 - t(v) \Rightarrow t(v) \leq \frac{1}{2} - \frac{b(v)}{2v} \Rightarrow (2\alpha - 1)v - \delta \leq 0 \quad (8)$$

$$\rho \frac{v+1}{3} \leq 1 - b(v) \Rightarrow b(v) \leq \frac{1}{2} - \frac{t(v)v}{2} \Rightarrow \alpha v^2 - (2\delta + 1)v + 3\beta \leq 0. \quad (9)$$

We choose $\alpha = \frac{1}{2} + \frac{1}{4}\delta$, which makes Restriction (8) tight. We choose $\beta = \delta$ which is the smallest value of β for which $b(v) \geq 0$ always. Substituting in Restriction (9) yields:

$$\left(\frac{1}{2}v^2 - v \right) + \delta \left(\frac{1}{4}v^2 - 2v + 3 \right) \leq 0. \quad (10)$$

Note that this constraint is tight at $v = 2$. The derivative of the left hand side is $v - 1 + \delta(v/2 - 2)$, which is negative at $v = 1$. Thus, so long as the constraint is satisfied at $v = 1$, it is satisfied for all $v \in [1, 2]$. For this, it suffices that $\frac{5}{4}\delta \leq \frac{1}{2}$, i.e. that $\delta \leq \frac{2}{5}$.

We choose $\delta = \frac{2}{5}$, and then $\beta = \frac{2}{5}$ and $\alpha = \frac{3}{5}$. Substituting in Equation (7) gives:

$$\rho = \frac{3 \left(\frac{3}{5}v^2 - \frac{2}{5}v + \frac{2}{5} \right)}{v^2 + v}. \quad (11)$$

It is a simple matter to check that the derivative of ρ is zero at $(2 + \sqrt{14})/5$ when $\rho \approx 0.89$.

Next, we argue that this is the best choice of parameters. Consider the choice of parameters $\alpha = \frac{1}{2} + \frac{1}{4}\delta + \alpha'$, where $\alpha' \leq 0$ (Restriction (8) at $v = 2$ forces $\alpha' \leq 0$), $\beta = \delta + \beta'$ (again, $b(2) \geq 0$ and $\gamma = 2\beta$ forces $\beta' \geq 0$), and $\delta = \frac{2}{5} + \delta'$ (again, $\delta' \leq 0$). Substituting into Restriction (9) with $v = 1$ yields:

$$\alpha' + 3\beta' - 2\delta' \leq 0. \quad (12)$$

Thus any change to the term $\alpha^2 - \delta v + \beta$ in ρ is negative: any increase due to β' must be offset by a decrease of at least $3\beta'v^2$ due to α' and any increase due to $-\delta'v$ is similarly offset by a decrease of $\frac{9}{4}\delta'v^2 + \delta'$ due to α' , α and β .

The above scheme handles the case that R_b is the middle bidder. Otherwise, w.l.o.g., suppose that the bottom bidder is R_b (if not, just switch the roles of the bottom and top items). Then for $v \leq 2$, the bottom bidder is assigned the bottom item, and the other two bidders each receive half the top item. For $2 \leq v \leq \sqrt{12}$, $t(v) = \frac{1}{4} - \frac{1}{v^2}$ and $b(v) = \frac{2}{v}$, and the other two bidders each receive $\frac{1}{4} + \frac{1}{v^2}$ of the top item. For $v > \sqrt{12}$, all three bidders receive one third of the top item. The worst approximation factor occurs at $v = \sqrt{12}$ and it has value $(12 - \sqrt{12})/11 \approx 0.77599$. □